

IN THE CLAIMS:

1. (Cancelled)
2. (Currently Amended) The method of claim 1 4, wherein the features are textual.
3. (Currently Amended) The method of claim 1 4, wherein the features are non-textual.
4. (Currently Amended) ~~The method of claim 1~~ A method for compressing an index file in an information retrieval system that retrieves information from a plurality of documents, each of the plurality of documents having features occurring therein, the method comprising the step of:
representing occurrence frequencies of the features in the plurality of documents in a compressed format in the index file, wherein the compressed format comprises a plurality of bin identifiers for a plurality of bins over which the occurrence frequencies are categorized, wherein
said representing step comprises the steps of:
mapping the occurrence frequencies into the plurality of bins; and
storing the bin identifiers in the index file, each of the bin identifiers identifying at least one of the bins to which at least one individual occurrence frequency is mapped.
5. (Previously Presented) The method of claim 4, further comprising the step of assigning a range of values of occurrence frequencies to each bin.
6. (Previously Presented) The method of claim 4, wherein the step of mapping results in one of the bins being empty .

7. (Previously Presented) The method of claim 5, wherein assigning comprises assigning a different range of values of occurrence frequencies, such that the different range represented by each of the plurality of bins contains a substantially same number of the occurrence frequencies.

8. (Original) The method of claim 4, wherein said mapping step respectively maps more than a single term and a corresponding occurrence frequency into each of the plurality of bins, the method further comprises the step of scoring at least one of the plurality of documents with respect to a query, and said scoring step comprises the step of computing an occurrence frequency for a given one of the plurality of bins as a weighted average of the occurrence frequencies contained within the given one of the plurality of bins.

9. (Original) The method of claim 4, wherein said mapping step respectively maps only a single term and a corresponding occurrence frequency into each of the plurality of bins, the method further comprises the step of scoring at least one of the plurality of documents with respect to a query, and said scoring step comprises the step of computing an occurrence frequency for a given one of the plurality of bins based on the single term and corresponding occurrence frequency mapped thereto.

10. (Original) The method of claim 4, further comprising the step of establishing bin boundaries for the plurality of bins based on a methodology employed to score the plurality of documents with respect to queries, the bin boundaries defining intervals within which the occurrence frequencies fall.

11. (Original) The method of claim 4, further comprising the steps of:
receiving a query having at least one term; and
computing a relevance score for at least one of the plurality of documents with respect to the query, based on the bin identifiers.

12. (Original) The method of claim 4, wherein the method is implemented by a program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform said method steps.

13. (Canceled)

14. (Currently Amended) The apparatus of claim ~~13~~ 16, wherein the features are textual.

15. (Currently Amended) The apparatus of claim ~~13~~ 16, wherein the features are non-textual.

16. (Currently Amended) ~~The apparatus of claim 13~~ An apparatus for compressing an index file in an information retrieval system that retrieves information from a plurality of documents, each of the plurality of documents having features occurring therein, the apparatus comprising:

a compression device for representing occurrence frequencies of the features in the plurality of documents in a compressed format in the index file, wherein the compressed format comprises a plurality of bin identifiers for a plurality of bins over which the occurrence frequencies are categorized, wherein said compression device comprises:

a bin generator for generating the plurality of bins and the corresponding plurality of bin identifiers, each of the plurality of bin identifiers respectively identifying one of the plurality of bins to which at least one individual occurrence frequency is mapped;

a mapping device for mapping the occurrence frequencies into the plurality of bins; and

a storage device for storing the bin identifiers in the index file.

17. (Previously Presented) The apparatus of claim 16, wherein said bin generator assigns a range of values of occurrence frequencies to each one of the bins.

18. (Previously Presented) The apparatus of claim 17, wherein said bin generator assigns a different range of values of occurrence frequencies, such that the different range represented by each of the plurality of bins contains a substantially same number of the occurrence frequencies.

19. (Original) The apparatus of claim 16, wherein said mapping device respectively maps more than a single term and a corresponding occurrence frequency into each of the plurality of bins, the apparatus further comprises a scoring device for scoring at least one of the plurality of documents with respect to a query by computing an occurrence frequency for a given one of the plurality of bins as a weighted average of the occurrence frequencies contained within the given one of the plurality of bins.

20. (Original) The apparatus of claim 16, wherein said mapping device respectively maps only a single term and a corresponding occurrence frequency into each of the plurality of bins, the apparatus further comprises a scoring device for scoring at least one of the plurality of documents with respect to a query by computing an occurrence frequency for a given one of the plurality of bins based on the single term and corresponding occurrence frequency mapped thereto.

21. (Original) The apparatus of claim 16, wherein said bin generator establishes bin boundaries for the plurality of bins based on a methodology employed to score the plurality of documents with respect to queries, the bin boundaries defining intervals within which the occurrence frequencies fall.

22. (Original) The apparatus of claim 16, further comprising a scoring device for computing a relevance score for at least one of the plurality of documents with respect to a query, based on the bin identifiers.

23. ~ 33. (Canceled)